



Proceedings of the  
**21st Annual Conference of  
the European Association  
for Machine Translation**

28–30 May 2018  
Universitat d'Alacant  
Alacant, Spain

*Edited by*

Juan Antonio Pérez-Ortiz  
Felipe Sánchez-Martínez  
Miquel Esplà-Gomis  
Maja Popović  
Celia Rico  
André Martins  
Joachim Van den Bogaert  
Mikel L. Forcada

*Organised by*



Universitat d'Alacant  
Universidad de Alicante

**transducens**  
research group





The papers published in this proceedings are —unless indicated otherwise— covered by the Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 International (CC-BY-ND 3.0). You may copy, distribute, and transmit the work, provided that you attribute it (authorship, proceedings, publisher) in the manner specified by the author(s) or licensor(s), and that you do not use it for commercial purposes. The full text of the licence may be found at <https://creativecommons.org/licenses/by-nc-nd/3.0/deed.en>.

© 2018 The authors

ISBN: 978-84-09-01901-4



# Contents

<b>Foreword from the General Chair</b> . . . . .	<b>v</b>
<b>Message from the Organising Committee Chair</b> . . . . .	<b>vii</b>
<b>Preface by the Programme Chairs</b> . . . . .	<b>ix</b>
<b>EAMT 2018 Committees</b> . . . . .	<b>xi</b>
<b>Sponsors</b> . . . . .	<b>xv</b>
<b>Invited Speech</b> . . . . .	<b>1</b>
<b>EAMT 2018 Best Thesis Award — Anthony C Clarke Award</b> . . . . .	<b>3</b>
Daniel Beck. <i>Gaussian Processes for Text Regression</i> . . . . .	5
<b>Special Feature</b> . . . . .	<b>7</b>
Barry Haddow. <i>The WMT Shared Tasks</i> . . . . .	9
<b>Research papers</b> . . . . .	<b>10</b>
Ruchit Agrawal, Marco Turchi and Matteo Negri. <i>Contextual Handling in Neural Machine Translation: Look behind, ahead and on both sides</i> . . . . .	11
Nora Aranberri and Jose A. Pascual. <i>Towards a post-editing recommendation system for Spanish–Basque machine translation</i> . . . . .	21
Duygu Ataman, Mattia Antonino Di Gangi and Marcello Federico. <i>Compositional Source Word Representations for Neural Machine Translation</i> . . . . .	31
Bogdan Babych. <i>Development and evaluation of phonological models for cognate identification</i> . . . . .	41
Sevilay Bayatli, Sefer Kurnaz, Ilnar Salimzyanov, Jonathan North Washington and Francis M. Tyers. <i>Rule-based machine translation from Kazakh to Turkish</i> . . . . .	49
Meriem Beloucif and Dekai Wu. <i>SRL for low resource languages isn’t needed for semantic SMT</i> . . . . .	59
Bram Bulté, Tom Vanallemeersch and Vincent Vandeghinste. <i>M3TRA: integrating TM and MT for professional translators</i> . . . . .	69
Sheila Castilho and Ana Guerberof Arenas. <i>Reading Comprehension of Machine Translation Output: What Makes for a Better Read?</i> . . . . .	79
Mara Chinae-Rios, Álvaro Peris and Francisco Casacuberta. <i>Are Automatic Metrics Robust and Reliable in Specific Machine Translation Tasks?</i> . . . . .	89
Mara Chinae-Rios, Germán Sanchis-Trilles and Francisco Casacuberta. <i>Creating the best development corpus for Statistical Machine Translation systems</i> . . . . .	99
Sandipan Dandapat and William Lewis. <i>Training Deployable General Domain MT for a Low Resource Language Pair: English–Bangla</i> . . . . .	109

Mattia A. Di Gangi and Marcello Federico. <i>Deep Neural Machine Translation with Weakly-Recurrent Units</i> . . . . .	119
Miguel Domingo and Francisco Casacuberta. <i>Spelling Normalization of Historical Documents by Using a Machine Translation Approach</i> . . . . .	129
Thierry Etchegoyhen, Eva Martínez Garcia, Andoni Azpeitia, Gorka Labaka, Iñaki Alegria, Itziar Cortes Etxabe, Amaia Jauregi Carrera, Igor Ellakuria Santos, Maite Martin and Eusebi Calonge. <i>Neural Machine Translation of Basque</i> . . . . .	139
M. Amin Farajian, Nicola Bertoldi, Matteo Negri, Marco Turchi and Marcello Federico. <i>Evaluation of Terminology Translation in Instance-Based Neural MT Adaptation</i>	149
Nisarg Jhaveri, Manish Gupta and Vasudeva Varma. <i>Translation Quality Estimation for Indian Languages</i> . . . . .	159
Tsz Kin Lam, Julia Kreutzer and Stefan Riezler. <i>A Reinforcement Learning Approach to Interactive-Predictive Neural Machine Translation</i> . . . . .	169
Jindřich Libovický, Thomas Brovelli Meyer and Bruno Cartoni. <i>Machine Translation Evaluation beyond the Sentence Level</i> . . . . .	179
Xutai Ma, Ke Li and Philipp Koehn. <i>An Analysis of Source Context Dependency in Neural Machine Translation</i> . . . . .	189
Mary Nurminen and Niko Papula. <i>Gist MT Users: A Snapshot of the Use and Users of One Online MT Tool</i> . . . . .	199
John E. Ortega, Weiyi Lu, Adam Meyers and Kyunghyun Cho. <i>Letting a Neural Network Decide Which Machine Translation System to Use for Black-Box Fuzzy-Match Repair</i> . . . . .	209
Zuzanna Parcheta, Germán Sanchis-Trilles and Francisco Casacuberta. <i>Data selection for NMT using Infrequent n-gram Recovery</i> . . . . .	219
Shantipriya Parida and Ondřej Bojar. <i>Translating Short Segments with NMT: A Case Study in English-to-Hindi</i> . . . . .	229
Alberto Poncelas, Gideon Maillette de Buy Wenniger and Andy Way. <i>Feature Decay Algorithms for Neural Machine Translation</i> . . . . .	239
Alberto Poncelas, Dimitar Shterionov, Andy Way, Gideon Maillette de Buy Wenniger and Peyman Passban. <i>Investigating Backtranslation in Neural Machine Translation</i>	249
Sander Tars and Mark Fishel. <i>Multi-Domain Neural Machine Translation</i> . . . . .	259
Vincent Vandeghinste, Lyan Verwimp, Joris Pelemans and Patrick Wambacq. <i>A Comparison of Different Punctuation Prediction Approaches in a Translation Context</i>	269
<b>User papers</b> . . . . .	<b>279</b>
Pierrette Bouillon, Sabrina Girletti, Paula Estrella, Jonathan Mutal, Martina Bellodi and Beatrice Bircher. <i>Integrating MT at Swiss Post’s Language Service: preliminary results</i> . . . . .	281
Sandipan Dandapat and Christian Federmann. <i>Iterative Data Augmentation for Neural Machine Translation: a Low Resource Case Study for English–Telugu</i> . . . . .	287
Morgan O’Brien. <i>Toward leveraging Gherkin Controlled Natural Language and Machine Translation for Global Product Information Development</i> . . . . .	293
Zuzanna Parcheta, Germán Sanchis-Trilles, Aliaksei Rudak and Siarhei Bratchenia. <i>Implementing a neural machine translation engine for mobile devices: the Lingvanex use case</i> . . . . .	297
Nicholas Ruiz, Srinivas Bangalore and John Chen. <i>Bootstrapping Multilingual Intent Models via Machine Translation for Dialog Automation</i> . . . . .	303
Tanja Schmidt and Lena Marg. <i>How to Move to Neural Machine Translation for Enterprise-Scale Programs—An Early Adoption Case Study</i> . . . . .	309

Nander Speerstra. <i>A Comparison of Statistical and Neural MT in a Multi-Product and Multilingual Software Company - User Study</i> . . . . .	315
<b>Translators' papers</b> . . . . .	<b>322</b>
Félix Do Carmo. <i>Does Machine Translation Really Produce Translations?</i> . . . . .	323
Laura Bruno, Antonio Miloro, Paula Estrella and Mariona Sabaté Carrove. <i>Pre-professional pre-conceptions?</i> . . . . .	325
Ariana López Pereira. <i>Determining translators' perception, productivity and post-editing effort when using SMT and NMT systems</i> . . . . .	327
Lorena Pérez Macías. <i>Machine translation post-editing in the professional translation market in Spain: a case study on the experience and opinion of professional translators</i> . . . . .	329
Pilar Sánchez-Gijón, Joss Moorkens and Andy Way. <i>Perception vs. Acceptability of TM and SMT Output: What do translators prefer?</i> . . . . .	331
Jeannette Stewart and Mikel L Forcada. <i>Learning to use machine translation on the Translation Commons Learn portal</i> . . . . .	333
Paloma Valenciano. <i>Use of NMT in Ubiquis Group</i> . . . . .	335
Anna Zaretskaya and Marcel Biller. <i>An In-house Translator's Experience with Machine Translation</i> . . . . .	337
<b>Project/product descriptions</b> . . . . .	<b>338</b>
Juan A. Alonso and Albert Llorens. <i>OctaveMT: Putting Three Birds into One Cage</i> . . . . .	339
Diego Bartolomé and José Masa. <i>TransPerfect's Private Neural Neural Machine Translation Portal</i> . . . . .	341
Giorgio Bernardinello. <i>Terminology validation for MT output</i> . . . . .	343
Nicola Bertoldi, Davide Caroselli and Marcello Federico. <i>The ModernMT Project</i> . . . . .	345
Pierrette Bouillon, Silvia Rodríguez Vázquez and Irene Strasly. <i>Developing a New Swiss Research Centre for Barrier-Free Communication</i> . . . . .	347
Itziar Cortes, Igor Leturia, Iñaki Alegria, Aitzol Astigarraga, Kepa Sarasola and Manex Garaio. <i>Massively multilingual accessible audioguides via cell phones</i> . . . . .	349
Thierry Etchegoyhen, Borja Anza Porras, Andoni Azpeitia, Eva Martínez Garcia, Paulo Vale, José Luis Fonseca, Teresa Lynn, Jane Dunne, Federico Gaspari, Andy Way, Victoria Arranz, Khalid Choukri, Vladimir Popescu, Pedro Neiva, Rui Neto, Maite Melero, David Perez Fernandez, Antonio Branco, Ruben Branco and Luis Gomes. <i>ELRI - European Language Resources Infrastructure</i> . . . . .	351
Ulrich Germann, Peggy van der Kreeft, Guntis Barzdins and Alexandra Birch. <i>The SUMMA Platform: Scalable Understanding of Multilingual Media</i> . . . . .	353
Judith Klein. <i>Smart Pre- and Post-Processing for STAR MT Translate</i> . . . . .	355
Samuel Läubli, Mathias Müller, Beat Horat and Martin Volk. <i>mtrain: A Convenience Tool for Machine Translation</i> . . . . .	357
Adrià Martín-Mor and Pilar Sánchez-Gijón. <i>Empowering Translators with MTradumàtica: A Do-It-Yourself statistical machine translation platform</i> . . . . .	359
Nicholas Ruiz, Andrew Ochoa, Jainam Shah, William Goethels and Sergio DelRio Diaz. <i>Speech Translation Systems as a Solution for a Wireless Earpiece</i> . . . . .	361
Lucia Specia. <i>Multi-modal Context Modelling for Machine Translation</i> . . . . .	363
Antonio Toral, Martijn Wieling, Sheila Castilho, Joss Moorkens and Andy Way. <i>Project PiPeNovel: Pilot on Post-editing Novels</i> . . . . .	365

Vincent Vandeghinste, Tom Vanallemeersch, Bram Bulté, Liesbeth Augustinus, Frank Van Eynde, Joris Pelemans, Lyan Verwimp, Patrick Wambacq, Geert Heyman, Marie-Francine Moens, Iulianna van der Lek-Ciudin, Frieda Steurs, Ayla Rigouts Terry, Els Lefever, Arda Tezcan, Lieve Macken, Sven Coppers, Jens Brulmans, Jan Van Den Bergh, Kris Luyten and Karin Coninx. <i>Smart Computer-Aided Translation Environment (SCATE): Highlights</i> . . . . .	367
Peggy van der Kreeft and Renars Liepins. <i>news.bridge – Automated Transcription and Translation for News</i> . . . . .	369
Eva Vanmassenhove and Christian Hardmeier. <i>Europarl Datasets with Demographic Speaker Information</i> . . . . .	371



# Foreword from the General Chair

As president of the European Association for Machine Translation (EAMT), it is a great pleasure for me to write the foreword to the Proceedings of the 21th annual conference of the EAMT.

The EAMT started organizing annual workshops in 1996; later, these workshops became annual conferences, and were hosted all around Europe. Years ago, the venue was steadily moving from west to east: from Barcelona (2009) to Saint-Raphaël (2010) to Leuven (2011) to Trento (2012) to Dubrovnik (2014) —after skipping one year to host the successful world-wide MT Summit 2013 in Nice—, but recently turned around to go west again at Antalya (2015), to go to Riga (2016), then Prague (2017) and now Alacant (2018). There will be no EAMT 2019, as it is the Association’s turn to organize the Machine Translation Summit, which will take place in Dublin, but EAMT 2020 will inevitably take place west from Alacant: it will be soon announced.

By the way, if you have not done so yet, and live in Europe, North Africa, or the Middle East, please consider joining the EAMT. Our membership rates are low, particularly for students and people not based in Europe. You will benefit from discounts when attending not only our conferences, but also the conferences held by our partner associations the Asia-Pacific Association for Machine Translation (AAMT) and the Association for Machine Translation in the Americas (AMTA). You will also have an exclusive chance to benefit from funding for your activities related to machine translation. And perhaps you can get even more involved and participate in serving the European machine translation community by becoming a member of the Executive Committee of the EAMT.

But let me go back to EAMT 2018. As in previous conferences, I am so happy to see the strong programme put together by our programme chairs: Maja Popović, research track chair, André Martins and Joachim van Bogaert, user track co-chairs, and Celia Rico, who will chair the new translators’ track, aiming at bringing machine translator researchers, developers, and vendors closer to the actual individuals using them. To accommodate this new track, EAMT 2018 will for the first time be a full three-day conference.

As in previous editions, there will also be a projects and products session showcasing the advance of machine translation in Europe. And, last but not least, I also feel very fortunate to have Sharon O’Brien from Dublin City University as our invited speaker.

EAMT 2018 would have never been possible without the generous offer to host and the hard work subsequently done by the local organizing committee at the Transducens research group of the Universitat d’Alacant, headed by Juan Antonio Pérez-Ortiz. I warmly thank my local colleagues (especially Juan Antonio, Felipe Sánchez-Martínez, and Miquel Esplà-Gomis) for putting EAMT 2018 together!

It is also with great pleasure that I thank our sponsors: Pangeanic (gold sponsor), Star Group and text&form (silver sponsors), Vicomtech (bronze sponsor), and Prompsit, Aper-tium, Linguaserve, and Unbabel (supporting sponsors), and ample support from the Universitat d’Alacant. Finally, I would like to thank EAMT 2018 attendees for coming to Alacant. I hope the conference leads to new friendships and fruitful collaboration.

Mikel L. Forcada  
EAMT President



# Message from the Organising Committee Chair

I want to take this opportunity to give you a big thank you for joining us in the 21st Annual Conference of the European Association for Machine Translation, EAMT 2018. On behalf of the organising committee, it is my pleasure to welcome you to Alacant. This year the Transducens research group and the Universitat d'Alacant proudly assume hosting the conference from the 28th to the 30th of May 2018. We decided to set the conference venue in downtown, but the main campus of our university is only six kilometres away. Regarded as one of the most beautiful European campuses, I encourage you to visit us there someday.

The city has held many different names: the Carthaginians called it Akra Leuka (white mountain), then the Romans changed its name to Lucentum, and the Moors—who ruled the region for a few centuries and started the building of the Santa Bàrbara castle on top of Mount Benacantil—called it Medina Laqant or al-Laqant. The Moorish name later resulted in the Catalan toponym Alacant and the Spanish Alicante. If you stand by the Postiguët beach and look up to the Mount Benacantil you will notice a rock formation that clearly resembles a man's face, a face that we have chosen as the main motif in the EAMT 2018 logo, where it is accompanied by some blue waves from the Mediterranean Sea. Legends tell us that the face is that of a Moorish king who was doomed to eternal damnation when her heartbroken daughter threw herself off the castle on to the rocks of Mount Benacantil after the king had disapproved her marriage with the one she truly loved. The king was condemned this way to watch all the lovers in the city and remember for all the eternity what he sadly forbade.

According to our predictions, this is going to be one of the most crowded editions ever of the EAMT conference. The unexpectedly high number of attendees have forced us to make some unavoidable last-minute changes that I hope will not negatively affect your enjoyment of the scientific and social activities of the conference.

We look forward to your active participation during the three days of the conference. Do not hesitate to ask questions when the session chairs invite you to do so. Please, contribute to make this edition of the conference a fruitful forum where a multidisciplinary group of researchers, developers, practitioners, leaders, vendors, users, and translators all share experiences and motivating ideas.

Finally, I would like to express my sincere appreciation to the persons and organisations that have made this conference possible: the European Association for Machine Translation, our gold sponsor (Pangeanic), silver sponsors (text&form and Star Group), bronze sponsor (Vicomtech), supporters (Apertium, Linguaserve, Prompsit, Unbabel), media sponsor (Multilingual), institutional partners (Universitat d'Alacant, Institut Universitari d'Investigació Informàtica), programme chairs (Maja Popović, Celia Rico, André Martins, Joachim Van den Bogaert), keynote speaker (Sharon O'Brien), and, finally but so importantly, my colleagues Miquel Esplà-Gomis, Mikel L. Forcada and Felipe Sánchez-Martínez who have worked extraordinarily hard to make your stay as pleasant and inspiring as possible.

Juan Antonio Pérez-Ortiz  
Universitat d'Alacant



# Preface by the Programme Chairs

It is our pleasure to welcome you to the 21st annual conference of the European Association for Machine Translation (EAMT) to be held in Alicante, Spain. We have really enjoyed serving as programme chairs for this edition of the conference. The EAMT conference has become the most important event in Europe in the area of machine translation for researchers, users and professional translators. This year, there are four different tracks: research, user and project/product track, as in previous editions, and for the first time, translators' – individual translators are invited to share their insights in the use of MT.

The research track concerns novel and significant research results in any aspect of machine translation and related areas while the user track reports users' experiences with machine translation in industry, government, NGOs, as well as innovative uses of MT. The project/product track offers project and products the opportunity to be presented to the wide audience of the conference. Finally, the machine translation community needs to hear the translators' voice in a fresh and unfiltered way and learn from their insights through the translators' track.

This year we have received 46 submissions to the research track, 16 submissions to the user track, 22 descriptions of projects and products and 10 submissions to the translators' track. Each submission to the research, user and translator tracks was peer reviewed by three independent members of the Programme Committee. In the research track, 27 papers (58.7%) were accepted for publication, whereas 7 papers (44%) were accepted for the user track and 8 submissions (80%) for the translators track. Aside from regular papers from the four tracks, the programme includes an invited talk by Sharon O'Brien from Dublin City University on "Human-centred translation technology". We will also have a presentation by Barry Haddow on "The WMT Shared Tasks", and a presentation by the winner of the EAMT Best Thesis Award.

We would like to thank the Programme Committee members whose names are listed below for their high quality reviews and recommendation which have been very useful for the Programme Chairs to make decisions. We would also like to thank all the authors for trying their best to incorporate the reviewers' suggestions when preparing the final versions of their papers. For the papers which were not accepted, we hope that the reviewers' comments will be useful for improving them. Special thanks to Mikel Forcada for all his help and advices.

Maja Popović Humboldt-Universität zu Berlin	Celia Rico Universidad Europea
--	-----------------------------------

André Martins Unbabel	Joachim Van den Bogaert CrossLang
--------------------------	--------------------------------------



# EAMT 2018 Committees

## General Chair

Mikel L. Forcada, EAMT President, Universitat d'Alacant

## Programme Chairs

### Research track

Maja Popović, Humboldt-Universität zu Berlin

### User track

André Martins, Unbabel

Joachim Van den Bogaert, CrossLang

### Translators' track

Celia Rico, Universidad Europea

## Organising committee

Juan Antonio Pérez-Ortiz, Universitat d'Alacant (chair)

Miquel Esplà-Gomis, Universitat d'Alacant

Felipe Sánchez-Martínez, Universitat d'Alacant

## Programme Committee

### Research track

Aleš Tamchyna, Memsource a.s.

Anabela Barreiro, INESC-ID

Andreas Eisele, European Commission, DGT

Andreas Guta, RWTH Aachen University

Andrei Popescu-Belis, IDIAP

Annette Rios Gonzales, University of Zurich

Antonio Toral, University of Groningen

Arianna Bisazza, University of Leiden

Bogdan Babych, University of Leeds

Carla Parra Escartín, Dublin City University

Carolina Scarton, The University of Sheffield

Christian Dugast, tech2biz  
Christian Federmann, Microsoft  
Christian Hardmeier , Uppsala University  
Clare Voss, ARL  
Constantin Orasan, University of Wolverhampton  
Cristina España i Bonet, UdS and DFKI  
Daniel Ortiz-Martínez, Universitat Politècnica de Valencia  
David Vilar Torres, Amazon  
Dimitar Shterionov, Dublin City University  
Ekaterina Lapshinova-Koltunski, Saarland University  
Eleftherios Avramidis, DFKI  
Eva Vanmassenhove, Dublin City University  
Evgeny Matusov, eBay  
Federico Gaspari, Dublin City University  
Felipe Sánchez-Martínez, Universitat d'Alacant  
Francisco Casacuberta, Universitat Politècnica de València  
Francisco Javier Guzman, Facebook  
Francis M. Tyers, Higher School of Economics  
Franck Burlot, LIMSI-CNRS  
François Yvon, LIMSI/CNRS et Université Paris-Sud  
George Foster, NRC  
Helena Caseli, Federal University of São Carlos (UFSCar)  
Houda Bouamor, Carnegie Mellon University  
Iacer Calixto, Dublin City University  
Irina Temnikova, University of Sofia  
Jan Niehues, Karlsruhe Institute of Technology  
Jerneja Žganec Gros, Alpineon R&D  
Joachim Daiber, Apple  
Joke Daems, Ghent University  
Jörg Tiedemann, University of Helsinki  
José G. C. de Souza, eBay Inc.  
Joss Moorkens, Dublin City University  
Juan Antonio Pérez-Ortiz, Universitat d'Alacant  
Laura Jehl, Universität Heidelberg  
Lieve Macken, Ghent University  
Luisa Bentivogli, FBK-irst  
Mārcis Pinnis, Tilde  
Marco Turchi, Fondazione Bruno Kessler  
Maria Nadejde, The University of Edinburgh  
Marianna Apidianaki, LIMSI-CNRS  
Marija Brkić, University of Rijeka  
Marion Weller-Di Marco, University of Amsterdam  
Mark Fishel, University of Tartu  
Markus Freitag, IBM  
Marta R. Costa-Jussà, Universitat Politècnica de Catalunya  
Martin Volk, University of Zurich  
Matteo Negri, Fondazione Bruno Kessler (FBK-irst)  
Matthias Huck, Ludwig Maximilian University of Munich  
Michel Simard, National Research Council Canada (NRC)



Mihaela Vela, Saarland University  
Miloš Stanojević, The University of Edinburgh  
Miquel Esplà-Gomis, Universitat d'Alacant  
Mireia Farrús, Universitat Pompeu Fabra  
Mirjam Sepesy Maučec, University of Maribor  
Nicola Ueffing, eBay  
Nizar Habash, New York University Abu Dhabi  
Núria Bel, Universitat Pompeu Fabra  
Parnia Bahar, RWTH Aachen University  
Philipp Koehn, Johns Hopkins University  
Philip Williams, The University of Edinburgh  
Preslav Nakov, Qatar Computing Research Institute  
Qun Liu, Dublin City University  
Rudolf Rosa, Charles University Prague  
Samuel Läubli, University of Zurich  
Sanja Štajner, University of Mannheim  
Sara Stymne, Uppsala University  
Sharon O'Brien, Dublin City University  
Sheila Castilho, Dublin City University  
Špela Vintar, University of Ljubljana  
Stephan Peitz, Apple  
Tamer Alkhoul, RWTH Aachen University  
Teresa Herrmann, Fujitsu  
Víctor M. Sánchez-Cartagena, Universitat d'Alacant  
Vincent Vandeghinste, Katholieke Universiteit Leuven  
Violeta Seretan, University of Geneva  
Yvette Graham, Dublin City University

## **User track**

Aljoscha Burchardt, DFKI  
Andrzej Zydrón, XTM Internation Ltd.  
Andy Way, ADAPT Centre — Dublin City University  
Arda Tezcan, Ghent University  
Arle Lommel, CSA Research  
Bianka Buschbeck, SAP  
Bruno Pouliquen, World Intellectual Property Organization  
Charlotte Tesselaar, LexisNexis  
Christine Bruckner, Freelance  
Fábio Kepler, Unbabel  
Gema Ramírez-Sánchez, Prompsit Language Engineering, S.L.  
Heidi Depraetere, CrossLang  
Heidi Van Hiel, Yamagata Europe  
Helena Moniz, Unbabel  
Joao Almeida Graca, Unbabel  
Jost Zetsche, IWG  
Julie Beliao, Unbabel  
Kim Harris, text & form  
Lena Marg, Welocalize  
Matiss Riktars, University of Latvia

Matthias Heyn, SDL  
Maxim Khalilov, Booking.com  
Miriam Kaeshammer, SAP  
Olga Beregovaya, Welocalize  
Ramon Astudillo, Unbabel / INESc-ID  
Samuel Läubli, University of Zürich  
Sara Szoc, CrossLang  
Tatjana Gornostaja, Tilde  
Teresa Herrmann, Fujitsu, Luxembourg  
Tony O’Dowd, Xcelerator Machine Translations Ltd.

### **Translators’ track**

Ana González, freelance translator  
Ana Guerberof, Adapt Centre, Dublin  
Enrique Torrejón, Deloitte, European Union Intellectual Property Office  
Gabriel Cabrera, freelance translator  
Ignacio Garcia, Western Sydney University  
Javier Mallo, freelance translator  
Javier Sánchez, Donnelley Language Solutions  
Julia Aymerich, PanAmerican Health Organization  
María Azqueta, Seprotec Multilingual Solutions  
Livia Florensa, CPSL  
Lucía Morado, Université de Genève  
Luis González, DGT – European Commission  
Manuel Mata, freelance translator  
Olga Blasco, Consultant, Business strategy  
Olga Torres-Hostench, Universitat Autònoma de Barcelona  
Pilar Sánchez-Gijón, Universitat Autònoma de Barcelona  
Rubén Rodríguez de la Fuente, PayPal  
Uwe Muegge, Anthrex  
Vanessa Enríquez Raído, The University of Auckland  
Vicenta Ten Soriano, SDL Trados Studio  
Willem Stoeller, Localization Institute

# Sponsors

Gold sponsor



Silver sponsors



## Bronze sponsors



## Supporters



## Media sponsor



[multilingual.com](http://multilingual.com)

## Institutional partners



# Invited Speech

## Human-centered translation technology

Sharon O'Brien, Dublin City University, Ireland

As AI drives advances in technology one recurring question is: what is the role of the human now and in the (near) future? This question is relevant for many disciplines, including medicine, law, accounting and, not least, translation. Translation is not a stranger to technology disruption and the modern translation pipeline is already highly technologised, at least in some sectors. However, there are benefits to focusing on users even in this high tech production pipeline. In my Keynote, I will suggest that we need to move from “computer-aided translation” and “human-in-the loop” to a human-centered translation technology (HCTT) paradigm. By focusing on three cohorts - professional translators, ad hoc translators, and end users - I will demonstrate how attention is shifting to HCTT and I will propose some research and development challenges for translation technology to embrace in order to position the human firmly in the centre of the design and use ecosystem.



# EAMT 2018 Best Thesis Award — Anthony C Clarke Award

Twelve PhD theses defended in 2017 were received as candidates for the 2018 edition of the EAMT Best Thesis Award, Anthony C Clarke Award, and all twelve were eligible. A panel of 41 reviewers was recruited to examine and score the theses, considering how challenging the problem tackled in each thesis was, how relevant the results are for machine translation as a field, and what the strength of its impact in terms of scientific publications was. It became very clear that 2017 was a very good year for PhD theses in machine translation. The scores of the best theses were very close, which made it very hard to select a winner. A panel of three EAMT Executive Committee members (Barry Haddow, Juan Antonio Pérez-Ortiz, and Mikel L. Forcada) was assembled to process the reviews and select a winner.

The panel has decided to grant the 2018 edition of the EAMT Best Thesis Award, Anthony C Clarke Award, to **Daniel Emilio Beck** for his thesis “Gaussian Processes for Text Regression”, University of Sheffield, supervised by Lucia Specia and Trevor Cohn.





# Gaussian Processes for Text Regression

Daniel Beck<sup>1</sup>

School of Computing and Information Systems  
The University of Melbourne, Australia  
d.beck@unimelb.edu.au

This thesis deals with the general problem of predicting numerical indicators from textual data. This task, which we call Text Regression, arises in a range of different applications in Natural Language Processing (NLP). For instance, in Quality Estimation (QE) (Blatz et al., 2004; Specia et al., 2009), sentences generated from Machine Translation (MT) systems are evaluated according to a task-based metric such as post-editing effort or time. In Emotion Analysis (EA) (Strapparava and Mihalcea, 2007), natural language sentences are assigned with numerical scores mapping the strength of a particular emotion (or a set of emotions).

Standard approaches for Text Regression rely on architectures similar to the ones used in classification tasks. These use engineered features and/or simple text representations such as bag-of-words (BOW), and make predictions in the form of single point estimates. These simplifying assumptions ignore important aspects of the data. Representations such as BOW ignore structural aspects of sentences and fails to capture structural linguistic phenomena such as word order. Point estimate predictions lack uncertainty information on the predicted variable, which can help subsequent decision making and is particularly important when annotations are noisy (such as post-editing time in QE).

The goal of this thesis is to advance the state-of-the-art in Text Regression by improving these two aspects: improved text representations and better uncertainty modelling in the response variables. In order to achieve that goal we propose

to use Gaussian Processes (GPs) (Rasmussen and Williams, 2006) as the regression model. GPs are a Bayesian kernelised framework which is considered the state-of-the-art in regression (Hensman et al., 2013). Perhaps surprisingly, GPs were not widely investigated in the context of NLP applications.<sup>2</sup> Therefore a secondary goal of this thesis is to disseminate GPs in the NLP community, in particular for regression tasks.

The theory behind Gaussian Processes regression makes it ideal to solve the two problems mentioned above. Since it models response variables as well-calibrated distributions, it naturally provides a measure of uncertainty over the predictions. Furthermore, by employing kernels as the underlying learning component, we can incorporate complex text representations through what we named *structural kernels*. Combining with the efficient model selection procedures provided by GPs, we show in this thesis how to essentially learn representations by enabling richer kernel parameterisations. In this thesis, we focus on string kernels (Lodhi et al., 2002; Cancedda et al., 2003) and tree kernels (Collins and Duffy, 2001; Moschitti, 2006) but the theory can easily be extended to other kinds of structures such as graph kernels (Vishwanathan et al., 2010).

We benchmark our approach in two Text Regression applications. The first one is Emotion Analysis, where we use a GP model with a soft string kernel using word embeddings for similarity calculation between words. We show that this proposed model can obtain better results compared to simpler baselines. For this task, we also propose a multi-task model which leverages multiple emotional labels and show how we can inspect GP

<sup>1</sup>This thesis was written while the author was a Ph.D. student at The University of Sheffield, United Kingdom.

© 2018 The authors. This article is licensed under a Creative Commons 3.0 licence, no derivative works, attribution, CC-BY-ND.

<sup>2</sup>Notable exceptions are Polajnar et al. (2011) and Cohn and Specia (2013).

hyperparameters to cluster similar emotions.

The second benchmark is Machine Translation Quality Estimation. In this task, we show that can obtain better results compared to baselines while also providing uncertainty estimates for predictions. More important, we show how to employ the predictive distributions in an asymmetric risk scenario, where over and underestimates of post-editing time have different costs. This is an example application where propagating full uncertainty information can be beneficial for further decision making in a translation pipeline. As another application example, we also show how to use uncertainty estimates to annotate QE datasets via active learning.

Finally, as mentioned before, this thesis also has the goal of disseminating Gaussian Processes among the NLP community. By providing the theoretical grounds and showcasing its application in two benchmarks, we hope that it will serve as a starting point for other NLP problems in the future.

Access to the full thesis is open and available at the White Rose eTheses repository ([etheses.whiterose.ac.uk/17619](http://etheses.whiterose.ac.uk/17619)).

## Acknowledgements

Daniel Beck was supported by funding from CNPq/Brazil (No. 237999/2012-9). The author would also like to thank his Ph.D. supervisors, Lucia Specia and Trevor Cohn, and his thesis examiners, Andreas Vlachos and Xavier Carreras.

## References

- Blatz, John, Erin Fitzgerald, George Foster, Simona Gandrabur, Cyril Goutte, Alex Kulesza, Alberto Sancho, and Nicola Ueffing. 2004. Confidence estimation for machine translation. In *Proceedings of the 20th Conference on Computational Linguistics*, pages 315–321.
- Cancedda, Nicola, Eric Gaussier, Cyril Goutte, and Jean-Michel Renders. 2003. Word-Sequence Kernels. *The Journal of Machine Learning Research*, 3:1059–1082.
- Cohn, Trevor and Lucia Specia. 2013. Modelling Annotator Bias with Multi-task Gaussian Processes: An Application to Machine Translation Quality Estimation. In *Proceedings of ACL*, pages 32–42.
- Collins, Michael and Nigel Duffy. 2001. Convolution Kernels for Natural Language. In *Proceedings of NIPS*, pages 625–632.
- Hensman, James, Nicolò Fusi, and Neil D. Lawrence. 2013. Gaussian Processes for Big Data. In *Proceedings of UAI*, pages 282–290.
- Lodhi, Huma, Craig Saunders, John Shawe-Taylor, Nello Cristianini, and Chris Watkins. 2002. Text Classification using String Kernels. *The Journal of Machine Learning Research*, 2:419–444.
- Moschitti, Alessandro. 2006. Making Tree Kernels practical for Natural Language Learning. In *EACL*, pages 113–120.
- Polajnar, Tamara, Simon Rogers, and Mark Girolami. 2011. Protein interaction detection in sentences via Gaussian Processes: a preliminary evaluation. *International Journal of Data Mining and Bioinformatics*, 5(1):52–72, jan.
- Rasmussen, Carl Edward and Christopher K. I. Williams. 2006. *Gaussian processes for machine learning*, volume 1. MIT Press Cambridge.
- Specia, Lucia, Nicola Cancedda, Marc Dymetman, Marco Turchi, and Nello Cristianini. 2009. Estimating the sentence-level quality of machine translation systems. In *Proceedings of EAMT*, pages 28–35.
- Strapparava, Carlo and Rada Mihalcea. 2007. SemEval-2007 Task 14 : Affective Text. In *Proceedings of SemEval*, pages 70–74.
- Vishwanathan, S. V. N., Nicol N. Schraudolph, Risi Kondor, and Karsten M. Borgwardt. 2010. Graph Kernels. *Journal of Machine Learning Research*, 11:1201–1242.

# Special Feature



# The WMT Shared Tasks

**Barry Haddow**

School of Informatics  
University of Edinburgh  
Edinburgh  
Scotland

bhaddow@staffmail.ed.ac.uk

## Abstract

The annual WMT Conference in Machine Translation has been running shared tasks since 2006. It started with a translation task based on Europarl, and has grown to include tasks on all aspects of MT corpus preparation, training and evaluation, including the flagship task on news translation. I will review the history of the task, lessons learnt, and plans for future tasks.

## 1 Introduction

We began organising shared tasks in machine translation at the Workshop in Machine Translation (WMT) in 2006, initially with a translation task based on Europarl. In later years, funding from the EU projects EuroMatrixPlus and MosesCore (FP7) and QT21 and Cracker (H2020), plus commercial sponsorship, enabled us to increase the number of tasks and to produce professionally translated, unseen test sets drawn from news texts for the translation task. In 2016 WMT became a conference (retaining the acronym) and in the last three years the number of shared tasks has varied between 7 and 10.

The shared tasks have covered translation (mainly news, but also other domains such as IT and biomedical and also more specialized tasks such as pronoun and multimodal), training (both tuning of SMT and training of NMT), reference-based evaluation, quality estimation, corpus preparation (document alignment and corpus cleaning) as well as automatic post-editing. The quality estimation task has included different subtasks on es-

timating the quality of MT output at word, sentence and document level, as well as trying to predict the post-editing effort required for a given MT output. The data from all the WMT tasks, including the training data, test data and task submissions is made available for future research and has been heavily used in academic publications.

In the news translation task we have tried to innovate in MT evaluation, whilst still providing for comparison with previous years. After several years using a *relative ranking* approach, where evaluators compare output from different systems, we switched to *direct assessment* (DA) in 2017. In DA, evaluators provide an assessment of adequacy on a scale from 0 to 100, which we find offers a reliable system ranking and a more interpretable and comparable final score. The news task covers a variety of languages, mainly European, with English–German and English–Czech as our “core” languages. We have included both low-resource (e.g. Estonian–English and Hindi–English) and high-resource (e.g. French–English) pairs, and we release our own parallel and monolingual data sets, as well as using standard sets like Europarl.

In this talk I will review the history of the tasks, the lessons learnt and plans for future tasks, focusing on the news translation task. I will explain how this task provides a common benchmark for comparing different MT systems, which helps to drive MT research. I will also show how running the task reveals difficulties and pitfalls in comparative evaluation of MT systems.

## 2 Website

The URL for the latest conference/task is [www.statmt.org/wmt18](http://www.statmt.org/wmt18), where you will find links to all previous conferences/workshops, tasks and papers.

© 2018 The authors. This article is licensed under a Creative Commons 3.0 licence, no derivative works, attribution, CC-BY-ND.

# Research papers