

WILKER FERREIRA AZIZ

Institute for Logic, Language and Computation
University of Amsterdam
Science Park 107, F2.11
Amsterdam, 1098 XG, Netherlands

w.aziz@uva.nl
<https://wilkeraziz.github.io>
<https://probabl1.github.io>
Tel: +31 (0) 631991009

Research

My research is concerned with enabling computers to understand and process human languages. To that end, I rely on results from probability theory, Bayesian statistics, optimisation, formal languages, and linguistics. I develop machine learning models as well as techniques to estimate such models efficiently from data. Because availability of supervision is a major bottleneck to statistical learning, my research has a strong component of unsupervised learning and other forms of latent variable modelling.

Selected Publications

- *Is MAP Decoding All You Need? The Inadequacy of the Mode in Neural Machine Translation.*
Bryan Eikema and Wilker Aziz. In *Coling*, 2020.
- *Efficient Marginalization of Discrete and Structured Latent Variables via Sparsity.*
Gonçalo M. Correia, Vlad Niculae, Wilker Aziz, and André F. T. Martins. In *NeurIPS*, 2020.
- *Effective Estimation of Deep Generative Language Models.*
Tom Pelsmaecker and Wilker Aziz. In *ACL*, 2020.
- *A Latent Morphology Model for Open-Vocabulary Neural Machine Translation.*
Duygu Ataman, Wilker Aziz, and Alexandra Birch. In *ICLR*, 2020.
- *Interpretable Neural Predictions with Differentiable Binary Variables.*
Jasmijn Bastings, Wilker Aziz and Ivan Titov. In *ACL*, 2019.
- *Block Neural Autoregressive Flow.*
Nicola De Cao, Wilker Aziz and Ivan Titov. In *UAI*, 2019.
- *Question Answering by Reasoning Across Documents with Graph Convolutional Networks.*
Nicola De Cao, Wilker Aziz and Ivan Titov. In *NAACL*, 2019.
- *A Stochastic Decoder for Neural Machine Translation.*
Philip Schulz, Wilker Aziz and Trevor Cohn. In *ACL*, 2018.
- *Fast collocation-based Bayesian HMM word alignment.*
Philip Schulz and Wilker Aziz. In *COLING*, 2016.
- *Exact Decoding for Phrase-Based Statistical Machine Translation.*
Wilker Aziz, Marc Dymetman and Lucia Specia. In *EMNLP*, 2014.

For a complete list refer to <https://wilkeraziz.github.io/publications.html>.

Education

2010–2014

Ph.D. Computational Linguistics

Research Institute in Information and Language Processing – University of Wolverhampton

Thesis: Exact Sampling and Optimisation in Statistical Machine Translation

Supervisors: Prof. Dr. Lucia Specia, Dr. Marc Dymetman, Prof. Dr. Ruslan Mitkov

Summary: I introduce an approach to exact optimisation and sampling based on a form of adaptive rejection sampling which addresses challenges in global optimisation and unbiased sampling in high-dimensional discrete spaces. In this view, an intractable goal distribution is upperbounded by a tractable proxy distribution which is then incrementally refined to be closer to the goal.

2005–2010

B.Sc. Computer Engineering

Escola de Engenharia de São Carlos - Universidade Estadual de São Paulo (USP)

Monograph: Lexical Substitution for Statistical Machine Translation

Summary: I propose a context model based on word co-occurrence and supervised learning to rank for cross-language lexical substitution.

Employment

- 01/2019–present **Assistant Professor, Institute for Logic, Language and Computation, Universiteit van Amsterdam, Netherlands**
Summary: My research group focuses on probabilistic models of language with a strong focus on unsupervised learning and latent variable modelling.
- 01/2015–12/2018 **Research Associate, Institute for Logic, Language and Computation, Universiteit van Amsterdam, Netherlands**
Summary: I joined the Statistical Language Processing and Learning Lab led by Professor Khalil Sima'an in January 2015 where I worked on several aspects of machine translation (e.g. word alignment, word reordering, and morphological analysis and generation) and paraphrasing employing log-linear, Bayesian, and deep generative models.
- 11/2013–12/2014 **Research Associate, Department of Computer Science, University of Sheffield, UK**
Summary: My work was funded by EPSRC under the MODIST (MOdelling DIscourse in Statistical Translation) project led by Prof. Dr. Lucia Specia. Discourse information typically requires nonlocal forms of parameterisation. I developed better decoding algorithms for SMT aiming at incorporating global features, particularly, I worked on a lazy incorporation of nonlocal parameterisation using a form of adaptive rejection sampling.
- 08/2013–12/2013 **Internship, Xerox Research Centre Europe (XRCE), Grenoble, France**
Summary: I worked with the Machine Learning for Document Access and Translation group under supervision of Dr. Marc Dymetman and Dr. Sriram Venkatapathy on developing an exact decoder/sampler for phrase-based SMT.
- 03/2009–02/2010 **Internship, Xerox Research Centre Europe (XRCE), Grenoble, France**
Summary: I worked with the Cross-Language Technologies group under supervision of Dr. Marc Dymetman and Dr. Lucia Specia on the use of context models and textual entailment to improve statistical machine translation coverage and quality.

Services

I often serve on the program committee of ML/CL/NLP/MT conferences:

- CL/NLP conferences: ACL, EMNLP, NAACL, Coling, IJCNLP, CoNLL
- ML conferences: NeurIPS (top reviewer 2019), ICML, ICLR
- MT conferences: WMT, EAMT, AMTA, MT Summit

and I review for CL/NLP/MT journals: Transactions of ACL (TACL), Natural Language Engineering, Computer Speech and Language, and Machine Translation. I serve (or have served) as area chair for *Machine Learning for Natural Language Processing* for *SEM 2019, ACL 2020, and ICLR 2021.

Patents

- U.S. Patent Application Filing: SAMPLING AND OPTIMIZATION IN PHRASED-BASED MACHINE TRANSLATION USING AN ENRICHED LANGUAGE MODEL REPRESENTATION
Inventor(s): Marc Dymetman; Wilker Aziz; Sriram Venkatapathy
U.S. Ser. No.: 13/750,338. Filed on: 01/25/2013
- U.S. Patent Application Filing: DYNAMIC BI-PHRASES FOR STATISTICAL MACHINE TRANSLATION
Inventor(s): Marc Dymetman; Wilker Aziz; Nicola Cancedda; Jean-Marc Coursimault; Vassilina Nikoulina; Lucia Specia.
U.S. Ser. No.: 12/780,040. Filed on: 05/20/2010

Teaching

I am highly committed to education and I believe in empowering students with knowledge in a healthy manner. Since 2015, I coordinate, design, and implement MSc and BSc courses offered at UvA covering topics such as:

- probabilistic graphical models and Bayesian methods for NLP
- approximate probabilistic inference: Markov chain Monte Carlo sampling and variational inference
- weighted automata and grammars, semirings, and deductive systems
- statistical and neural approaches to natural language processing

I mostly design my own materials and value depth and quality.

Course	<i>Deep Learning for Natural Language Processing</i>
Role	Lecturer (2019–present) offered in collaboration with Christof Monz (IvI)
Programme	Master's of AI (UvA)
Description	The course covers advanced supervised and unsupervised learning techniques in natural language processing with a focus on statistical learning powered by deep neural networks.

Course	<i>Basic Probability Theory</i>
Role	Coordinator (2019–present)
Programme	Master's of Logic (UvA)
URL	https://basicprobability.github.io
Description	The course covers the basics of combinatorics, axiomatic probability theory, discrete and continuous random variables, and maximum likelihood estimation.

Course	<i>Natural Language Models and Interfaces</i>
Role	Coordinator (2018–present)
Programme	Bachelor's of AI (UvA)
URL	https://uva-slpl.github.io/nlmi/
Description	The course covers some of the essential techniques in natural language processing with a focus on language modelling and word representation.

Course	<i>Unsupervised Language Learning</i>
Role	Lecturer (2018) offered in collaboration with Ekaterina Shutova (ILLC)
Programme	Master's of AI (UvA)
URL	https://uva-slpl.github.io/ull/
Description	The course covers advanced unsupervised learning techniques in natural language processing with a focus on meaning representation (including deep generative models of word and sentence representation).

Course	<i>Natural Language Processing 2</i>
Role	Lecturer (2015-2017)
Programme	Master's of AI (UvA)
URL	https://uva-slpl.github.io/nlp2/
Description	The course covers structure prediction problems related to translation (e.g. unsupervised alignment, synchronous grammar induction, statistical and neural MT).

For a complete list of courses and projects including available material and project outcomes, please refer to <https://wilkeraziz.github.io/teaching.html>. For invited lectures and talks, please refer to <https://wilkeraziz.github.io/talks.html>.

I have recently developed a tutorial on variational inference and deep generative models for NLP audiences. This is a joint effort with Philip Schulz, and we have been taking this tutorial (or parts of it) to diverse audiences at universities (e.g. Melbourne, Monash, Macquarie, Amsterdam, Heidelberg, and Instituto Superior Técnico in Lisbon), companies (e.g. Amazon, Naver Labs, Yandex), and international conferences (ACL 2018). For more information check our schedule and available resources: <https://vitutorial.github.io>.

Supervision

I see supervision as a very special form of training where the goal is to transfer not only skills, but also an attitude towards research and towards people. Above all, I respect my students' personal lives and incentivise them to do the same. My experience with one-to-one supervision includes PhD, MSc, and BSc students.

	ONGOING	ROLE
PHD	Bryan Eikema (2019–present; UvA) <i>Data-Efficient Probabilistic Neural Machine Translation</i>	supervisor
	Nicola De Cao (2019–present; UvA) <i>Role of Knowledge Bases and Logical Reasoning in Question Answering</i>	co-supervisor
	COMPLETE	ROLE
PHD	Jasmijn Bastings (2017–2020; UvA) <i>A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing</i>	co-supervisor
	Philip Schulz (2015–2018; UvA) <i>Latent Variable Models for Machine Translation and How to Learn Them</i>	co-supervisor
	Joachim Daiber (2016–2018; UvA) <i>Typologically Robust Statistical Machine Translation</i>	co-supervisor
	Miloš Stanojević (2015–2017; UvA) <i>Permutation Forests for Modeling Word Order in Machine Translation</i>	co-supervisor
MSC	Lina Murady (2019; UvA) <i>Probabilistic Models for Joint Classification and Rationale Extraction</i>	supervisor
	Mathijs Pieters (2019; UvA) <i>Transparency for Text Classification Models</i>	supervisor
	Eelco van der Wel (2019; UvA) <i>Improving Controllable Generation with Semi-Supervised Deep Generated Models</i>	supervisor
	Ruben Gerritse (2019; UvA) <i>Co-Training Generative Neural Machine Translation Models</i>	supervisor
	Nuno Mota (2018–2019; UvA) <i>Textual (Generalised) Any-Shot Learning - The Case of Relation Classification</i>	supervisor
	Akash Raj (2019; UvA) <i>Semi-supervised Morphological Reinflection using Rectified Random Variables</i>	supervisor
	Dhruba (2019; UvA) <i>Supervised Neural Disease Normalization</i>	supervisor
	Daan van Stigt (2018–2019; UvA) <i>Neural Language Models with Latent Syntax</i>	supervisor
	Tom Pelsmaeker (2017–2018; UvA) <i>Effective Estimation of Deep Generative Models of Language</i>	supervisor
	Bryan Eikema (2017–2018; UvA) <i>Semi-Supervised Learning for Neural Machine Translation</i>	supervisor
Sanders Bijl de Vroe (2017; UU) <i>Character-level Neural Architectures for Jointly Predicting Word Alignments and Word-internal Structure in Morphologically Complex Languages</i>	supervisor	

For more information and links to theses, please refer to <https://wilkeraziz.github.io/people.html>.

Personal

I am a native Portuguese speaker who most of the time expresses himself in English (fluent) or Italian (intermediate), and who used to know some French. I play the guitar and try to keep Brazilian bossa alive, though on a very small scale. I love pets (all of them!), but I find dogs objectively superior to cats.